

KATHMANDU UNIVERSITY
End Semester Examination
May/June, 2022

Marks Scored:

Level : B.E./B.Sc.
Year : II

Course : COMP 232
Semester : II

Exam Roll No. :

Time: 30 mins.

F. M. : 10

Registration No.:

Date :

SECTION "A"

[20Q. × 0.5 = 10 marks]

Encircle the most appropriate answer.

- Which of the following statements is not **CORRECT**?
 - A database is a collection of related data necessary to manage an organization.
 - A DBMS is a software system to facilitate the creation and maintenance of a computerized database.
 - Databases may be more expensive to maintain compared to files because of backup and recovery needs.
 - A physical view represents how the users view the data.
- Which of the following is **NOT** an advantages of database management systems?
 - Increases redundancy
 - Restricts unauthorized access
 - Provides backup and recovery
 - Enforces integrity constraints
- If a student can attend 5 courses, and different professors can offer the same courses, what type of relationship exists between students and professors?
 - many-to-many
 - one-to-many
 - one-to-one
 - many-to-one
- A derived attribute
 - must be stored physically within in the database.
 - need not be physically stored within the database.
 - has many values.
 - is represented by a diamond shape in an ER diagram.
- Which of the following statements best describes the function of an entity relation model?
 - An ER model is concerned primarily with a logical view of the data and secondly with the physical implementation.
 - An ER model is concerned primarily with a physical implementation and secondly with the logical view of the data.
 - An ER model provides a view of the logic of the data and not the physical implementation.
 - An ER model is entirely concerned with the modeling the physical implementation.
- Which of the following operations results in a relation that includes all tuples that are in both R and S?
 - join
 - minus
 - union
 - intersection
- If R has x tuples, and S has y tuples, then how many tuples will $R \times S$ have?
 - $x + y$
 - $x * y$
 - It depends on how many tuples have the same values for the common attributes.
 - 0 if there is no common attributes.

8. Which of the following is **NOT** an aggregate function in SQL?
 a. SUM b. COUNT c. AVERAGE d. LARGEST
9. Which of the following query is equivalent to the one below:
 SELECT item_no FROM item WHERE price > 1000 AND price <1500;
 a. SELECT * FROM item WHERE price > 1000 AND price <1500;
 b. SELECT item_no FROM item WHERE price IN (1000, 1500);
 c. SELECT item_no FROM item WHERE price BETWEEN 101 AND 1499;
 d. SELECT item_no FROM item WHERE price BETWEEN 100 AND 1500;
10. Which of the following is the special operator used to check whether a subquery returns any rows?
 a. EXISTS b. BETWEEN c. IN d. LIKE
11. Which of the following statements is **FALSE**?
 a. BCNF is stricter than 3NF.
 b. Any relation with two singled valued attributes is in BCNF.
 c. Every relation in BCNF is also in 3NF.
 d. No relation can be in both BCNF and 3NF.
12. Consider a relation R(A, B, C, D, E, F) with the following set of functional dependencies: F= {A → B, AD → C, B → F, A → E}. Which of the following is a candidate key?
 a. A b. ACD c. AD d. AC
13. Given the following relation state, which of the following functional dependencies may be a valid one?

A	B	C	D
A1	B1	C1	D1
A1	B2	C1	D2
A2	B2	C2	D3
A2	B3	C2	D4
A3	B3	C2	D3

- a. A → C b. C → A c. B → C d. C → B
14. A transaction is a _____ unit of work that must be either entirely completed or aborted.
 a. timed b. practical c. logical d. physical
15. Which of the following does hold in a consistent database?
 a. All tables have foreign keys.
 b. All tables are normalized.
 c. All data integrity constraints are satisfied.
 d. All SQL statements update only one table at a time.
16. Which of the following is **CORRECT**?
 a. The wait-for graph is used to detect conflict serializability.
 b. There is no deadlock in the timestamp-based concurrency control.
 c. The two-phase locking protocol ensures freedom from deadlocks.
 d. Strict two-phase locking is difficult to implement and most database systems do not support it.

17. As long as two transactions, T1 and T2, access _____ data, there is no conflict, and the order of execution is irrelevant to the final outcome.
- a. shared b. common c. locked d. unrelated
18. What does the write-ahead logging (WAL) protocol mean?
- a. The log record for an operation should be written before the actual data is written.
b. All log records should be written before a new transaction begins.
c. All log records should be kept in a separate database.
d. The writing of data item should be done ahead of any logging operations.
19. Which of the following statements is **TRUE**?
- a. A checkpoint checks for deadlock situations.
b. A checkpoint checks for non-serializable schedules.
c. A checkpoint is an entry in a log file.
d. A checkpoint is needed to ensure system privacy.
20. NoSQL databases store data in different formats? Which of the format is **NOT** associated with NoSQL databases?
- a. Key-value pair b. Graph c. Document d. Table



KATHMANDU UNIVERISTY
End Semester Examination
May/June, 2022

Level : B.E./B.Sc.
Year : II
Time : 2 hrs. 30 mins.

Course : COMP 232
Semester : II
F.M. : 40

SECTION "B"

[6Q. × 4 = 24 marks]

Attempt *ANY SIX* questions.

1. What are the major components of a database management system?
2. Suppose you need to design a database for a website that allows users to book flights. Users must be able to book multiple flights from various airlines but are not allowed to book multiple flights with the same source and destination airports for the same day. Draw an ER diagram for this system.
3. Define functional dependency. Compare and contrast full, partial and transitive dependencies. [1 + 3]
4. What problems can arise when more than one transaction are executed concurrently? In spite of those problems, why is concurrent execution of transactions preferred? [3 + 1]
5. How do log-based recovery mechanisms ensure atomicity despite failures?
6. What do you understand by ACID-compliant databases? What are the benefits of NoSQL and NewSQL databases over SQL databases? [2 + 2]
7. Write short notes on (*ANY TWO*): [2 + 2]
 - a. Different types of relationships
 - b. Integrity constraints
 - c. Transaction and its properties

SECTION "C"

[2Q. × 8 = 16 marks]

Attempt *ANY TWO* questions.

8. Consider the following tables:
item(id, name, price, category_id, is_available)
category(id, name, parent_category_id)
order(id, date, total_amount)
order_details(order_id, item_id, quantity, amount)
 - i. Write an SQL script for creating category table. parent_category_id is a foreign key that references category table. [2]
 - ii. Give an expression in relational algebra for the following queries: [2]
 - a. Find all items in the category 'clothes'.
 - iii. Give an expression in SQL for the following queries: [2]
 - a. Find today's most ordered items.
 - b. There is a 10% discount on all clothes but not on shoes. List the name, original price, and discounted price, if any, of all items in the categories 'clothes' and 'shoes'. [2]

9. Consider a relation schema $r(A, B, C, D, E, F, G)$ with the following functional dependencies: $F = \{AB \rightarrow CD, B \rightarrow CF, A \rightarrow F, E \rightarrow A, G \rightarrow E\}$.

- i. List all candidate keys. [2]
- ii. Determine the Normal Form of this relation. [2]
- iii. If r is not in BCNF, decompose it into a collection of BCNF relations. [2]
- iv. What are the advantages and disadvantages of normalization? [2]

10. Consider the following schedule:

T1	T2	T3	T4
Read(Y)	Read(X)		
Write(Y)	Write(X)	Read(X)	Read(Y)
		Write(X)	Read(X)
	Read(Y)		Write(Y)
	Write(Y)		

- i. Is this schedule conflict-serializable? [2]
- ii. If lock and unlock instructions are added as below, will there be a deadlock? [2]

T1	T2	T3	T4
Lock-X(Y)			
Read(Y)	Lock-X(X)		
	Read(X)		
Write(Y)	Write(X)	Lock-X(X)	Lock-X(Y)
Unlock(Y)		Read(X)	Read(Y)
			Lock-S(X)
		Write(X)	Read(X)
		Unlock(X)	
	Lock-X(Y)		Write(Y)
	Read(Y)		Unlock(Y)
	Unlock(X)		Unlock(X)
	Write(Y)		
	Unlock(Y)		

- iii. Explain one strategy to resolve a deadlock. Explain one strategy to prevent a deadlock. [1 + 1]
- iv. Do you think preventing a deadlock is more important than detecting and then recovering from a deadlock? Justify your answer. [2]